


2

1992
Executive Research Project
S33

High Performance Computing and Communications Initiative "A Paradigm for National Industrial Policy?"

Lieutenant Colonel
Marlin G. Forbes, Jr.
U.S. Army

Faculty Research Advisor
Dr. Linda S. Brandt



The Industrial College of the Armed Forces
National Defense University
Fort McNair, Washington, D.C. 20319-6000



Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS	
2. SECURITY CLASSIFICATION AUTHORITY N/A		3. DISTRIBUTION / AVAILABILITY OF REPORT Distribution Statement A: Approved for public release; distribution is unlimited.	
5. DECLASSIFICATION / DOWNGRADING SCHEDULE N/A		5. MONITORING ORGANIZATION REPORT NUMBER(S) Same	
PERFORMING ORGANIZATION REPORT NUMBER(S) NDU-ICAF-92- 333		7a. NAME OF MONITORING ORGANIZATION National Defense University	
a. NAME OF PERFORMING ORGANIZATION Industrial College of the Armed Forces		6b. OFFICE SYMBOL (If applicable) ICAF-FAP	
c. ADDRESS (City, State, and ZIP Code) Fort Lesley J. McNair Washington, D.C. 20319-6000		7b. ADDRESS (City, State, and ZIP Code) Fort Lesley J. McNair Washington, D.C. 20319-6000	
a. NAME OF FUNDING / SPONSORING ORGANIZATION		8b. OFFICE SYMBOL (If applicable)	
9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER		10. SOURCE OF FUNDING NUMBERS	
b. ADDRESS (City, State, and ZIP Code)		PROGRAM ELEMENT NO.	
		PROJECT NO.	
		TASK NO.	
		WORK UNIT ACCESSION NO.	
1. TITLE (Include Security Classification) <i>High Performance Computing and Communications Initiative "A Paradigm for Future Industrial Policy?"</i>			
12. PERSONAL AUTHOR(S) <i>Marlin J. Forbes, Jr.</i>			
13a. TYPE OF REPORT Research		13b. TIME COVERED FROM Aug 91 TO Apr 92	
14. DATE OF REPORT (Year, Month, Day) April 92		15. PAGE COUNT 51	
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) SEE ATTACHED			
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Judy Clark		22b. TELEPHONE (Include Area Code) (202) 475-1889	
		22c. OFFICE SYMBOL ICAF-FAP	

EXECUTIVE SUMMARY

The United States finds itself at a very important strategic juncture. Technological leadership, once the hallmark of American industrial might, is continuing to erode, and with it industrial productivity and personal standards of living. The success of America's trading partners, particularly Japan, coupled with the potential threat posed by a united Europe, has heightened interest in the formulation of a strategic, long-term, and highly focused national industrial policy.

In that regard, the High-Performance Computing Act and the High-Performance Computing and Communications Initiative, while directly responding to threats to technological leadership, indirectly offers a paradigm for the formulation and execution of industrial policy. This research specifically looks at the high-performance domestic and international computer industry, addresses the requirements for continued government research and development spending, offers specific recommendations for improving government and industry cooperative efforts (consortia) and examines how future government policy could be better focused to improve productivity and increase national wealth.

1992
Executive Research Project
S33

High Performance Computing and Communications Initiative "A Paradigm for National Industrial Policy?"

Lieutenant Colonel
Marlin G. Forbes, Jr.
U.S. Army

Faculty Research Advisor
Dr. Linda S. Brandt



The Industrial College of the Armed Forces
National Defense University
Fort McNair, Washington, D.C. 20319-6000

DISCLAIMER

This research report represents the views of the author and does not necessarily reflect the official opinion of the Industrial College of the Armed Forces, the National Defense University, or the Department of Defense.

This document is the property of the United States Government and is not to be reproduced in whole or in part for distribution outside the federal executive branch without permission of the Director of Research and Publications, Industrial College of the Armed Forces, Fort Lesley J. McNair, Washington, D.C. 20319-6000.

1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019, 2020, 2021, 2022, 2023, 2024, 2025, 2026, 2027, 2028, 2029, 2030, 2031, 2032, 2033, 2034, 2035, 2036, 2037, 2038, 2039, 2040, 2041, 2042, 2043, 2044, 2045, 2046, 2047, 2048, 2049, 2050, 2051, 2052, 2053, 2054, 2055, 2056, 2057, 2058, 2059, 2060, 2061, 2062, 2063, 2064, 2065, 2066, 2067, 2068, 2069, 2070, 2071, 2072, 2073, 2074, 2075, 2076, 2077, 2078, 2079, 2080, 2081, 2082, 2083, 2084, 2085, 2086, 2087, 2088, 2089, 2090, 2091, 2092, 2093, 2094, 2095, 2096, 2097, 2098, 2099, 2100, 2101, 2102, 2103, 2104, 2105, 2106, 2107, 2108, 2109, 2110, 2111, 2112, 2113, 2114, 2115, 2116, 2117, 2118, 2119, 2120, 2121, 2122, 2123, 2124, 2125, 2126, 2127, 2128, 2129, 2130, 2131, 2132, 2133, 2134, 2135, 2136, 2137, 2138, 2139, 2140, 2141, 2142, 2143, 2144, 2145, 2146, 2147, 2148, 2149, 2150, 2151, 2152, 2153, 2154, 2155, 2156, 2157, 2158, 2159, 2160, 2161, 2162, 2163, 2164, 2165, 2166, 2167, 2168, 2169, 2170, 2171, 2172, 2173, 2174, 2175, 2176, 2177, 2178, 2179, 2180, 2181, 2182, 2183, 2184, 2185, 2186, 2187, 2188, 2189, 2190, 2191, 2192, 2193, 2194, 2195, 2196, 2197, 2198, 2199, 2200, 2201, 2202, 2203, 2204, 2205, 2206, 2207, 2208, 2209, 2210, 2211, 2212, 2213, 2214, 2215, 2216, 2217, 2218, 2219, 2220, 2221, 2222, 2223, 2224, 2225, 2226, 2227, 2228, 2229, 2230, 2231, 2232, 2233, 2234, 2235, 2236, 2237, 2238, 2239, 2240, 2241, 2242, 2243, 2244, 2245, 2246, 2247, 2248, 2249, 2250, 2251, 2252, 2253, 2254, 2255, 2256, 2257, 2258, 2259, 2260, 2261, 2262, 2263, 2264, 2265, 2266, 2267, 2268, 2269, 2270, 2271, 2272, 2273, 2274, 2275, 2276, 2277, 2278, 2279, 2280, 2281, 2282, 2283, 2284, 2285, 2286, 2287, 2288, 2289, 2290, 2291, 2292, 2293, 2294, 2295, 2296, 2297, 2298, 2299, 2300, 2301, 2302, 2303, 2304, 2305, 2306, 2307, 2308, 2309, 2310, 2311, 2312, 2313, 2314, 2315, 2316, 2317, 2318, 2319, 2320, 2321, 2322, 2323, 2324, 2325, 2326, 2327, 2328, 2329, 2330, 2331, 2332, 2333, 2334, 2335, 2336, 2337, 2338, 2339, 2340, 2341, 2342, 2343, 2344, 2345, 2346, 2347, 2348, 2349, 2350, 2351, 2352, 2353, 2354, 2355, 2356, 2357, 2358, 2359, 2360, 2361, 2362, 2363, 2364, 2365, 2366, 2367, 2368, 2369, 2370, 2371, 2372, 2373, 2374, 2375, 2376, 2377, 2378, 2379, 2380, 2381, 2382, 2383, 2384, 2385, 2386, 2387, 2388, 2389, 2390, 2391, 2392, 2393, 2394, 2395, 2396, 2397, 2398, 2399, 2400, 2401, 2402, 2403, 2404, 2405, 2406, 2407, 2408, 2409, 2410, 2411, 2412, 2413, 2414, 2415, 2416, 2417, 2418, 2419, 2420, 2421, 2422, 2423, 2424, 2425, 2426, 2427, 2428, 2429, 2430, 2431, 2432, 2433, 2434, 2435, 2436, 2437, 2438, 2439, 2440, 2441, 2442, 2443, 2444, 2445, 2446, 2447, 2448, 2449, 2450, 2451, 2452, 2453, 2454, 2455, 2456, 2457, 2458, 2459, 2460, 2461, 2462, 2463, 2464, 2465, 2466, 2467, 2468, 2469, 2470, 2471, 2472, 2473, 2474, 2475, 2476, 2477, 2478, 2479, 2480, 2481, 2482, 2483, 2484, 2485, 2486, 2487, 2488, 2489, 2490, 2491, 2492, 2493, 2494, 2495, 2496, 2497, 2498, 2499, 2500, 2501, 2502, 2503, 2504, 2505, 2506, 2507, 2508, 2509, 2510, 2511, 2512, 2513, 2514, 2515, 2516, 2517, 2518, 2519, 2520, 2521, 2522, 2523, 2524, 2525, 2526, 2527, 2528, 2529, 2530, 2531, 2532, 2533, 2534, 2535, 2536, 2537, 2538, 2539, 2540, 2541, 2542, 2543, 2544, 2545, 2546, 2547, 2548, 2549, 2550, 2551, 2552, 2553, 2554, 2555, 2556, 2557, 2558, 2559, 2560, 2561, 2562, 2563, 2564, 2565, 2566, 2567, 2568, 2569, 2570, 2571, 2572, 2573, 2574, 2575, 2576, 2577, 2578, 2579, 2580, 2581, 2582, 2583, 2584, 2585, 2586, 2587, 2588, 2589, 2590, 2591, 2592, 2593, 2594, 2595, 2596, 2597, 2598, 2599, 2600, 2601, 2602, 2603, 2604, 2605, 2606, 2607, 2608, 2609, 2610, 2611, 2612, 2613, 2614, 2615, 2616, 2617, 2618, 2619, 2620, 2621, 2622, 2623, 2624, 2625, 2626, 2627, 2628, 2629, 2630, 2631, 2632, 2633, 2634, 2635, 2636, 2637, 2638, 2639, 2640, 2641, 2642, 2643, 2644, 2645, 2646, 2647, 2648, 2649, 2650, 2651, 2652, 2653, 2654, 2655, 2656, 2657, 2658, 2659, 2660, 2661, 2662, 2663, 2664, 2665, 2666, 2667, 2668, 2669, 2670, 2671, 2672, 2673, 2674, 2675, 2676, 2677, 2678, 2679, 26

Page

DTIC QUALITY INSPECTED

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Special

A-1

HPCCI & Industrial Policy

Endnotes	38
Bibliography	40

INTRODUCTION

Within the past ten years, the United States, Japan and the European Community have embarked upon massively funded research and development programs in computers, information processing (read software and system integration), and communications technology. These initiatives are based upon a simple assumption, "the nation that dominates this information-processing field will possess the keys to world leadership in the twenty-first century."¹ This is especially true of the area of information processing technology known as high performance or super computing. The purpose of this research is to first gain an appreciation of the importance of high performance computing, look at the present state of the industry internationally and domestically, examine current domestic policy and determine if this policy could be used as a paradigm for national industrial policy. First, I want to examine what a super or high performance computer* is and why it is important to national security, from both an economic and defense point of view.

What is a super computer?

Obviously, speed is what makes the super computer "super". It provides for the first time, an opportunity for scientists to model events in minutes that previously took hours or days of computer processing time. The processing speed of computers is measured in millions of instructions per second or MIPS. For purposes of comparison, it is generally agreed that computers can be roughly divided into five categories based upon the speed at which they operate²;

* for the purpose of this research, the terms "super" or "high speed" computing will be used interchangeably

Computer Categories

1. Personal computers or Microprocessors	- 0.5 MIPS
2. Minicomputers	0.5 - 1.5 MIPS
3. Supermini computers	1.0 - 5.0 MIPS
4. Mainframe computers	8.0 - 20.0 MIPS
5. Super computers	50.0 - MIPS

While frequent changes of technology continue to blur the distinctions and gaps between categories and leave room for ambiguity, clearly super computers represent the upscale end of the computing industry. For our purposes, they are *the most powerful general purpose computers available for large scale scientific computation at any given time*. Since these computers are usually reserved for complex scientific calculations on floating point numbers, their speed is measured not in MIPS, but in millions, billions and trillions of floating point operations per second, or megaflops, gigaflops or teraflops, respectively. As technology continues to improve, we can expect continued shifts in these categories.

While computing in general has left the confines of laboratories and universities, the super computer epitomizes this transition. Recently, super computing has been thrust into the United States policy arena and represents a potential paradigm for government and private cooperation to ensure technical superiority in research and development, through implementation of super computing consortia. In particular, the last two years represent a watershed in government policy toward industry, particularly in super computing research and development. This change in policy could presage the way government will work with industry in other critical technology areas.

What are the issues? What is the U.S. doing about them? What programs are underway in Japan and the European Community? Does

HPCCI & Industrial Policy

more needs to be done? What lessons have been learned that could be used to shape government policy in high-risk, expensive and critical technology areas? The purpose of this paper will be to address these issues, with particular emphasis on government policy, domestic and foreign, in the development and use of super computers.

State of the Industry

At the present time Cray is the only major long-term, commercially successful manufacturer of super computers in the United States, although there are a number of startup companies. Control Data Corporation (CDC) stopped production of a line of super computers in 1989 due to continuing financial losses and domination of the marketplace by Cray. The Japanese have also shown determination to penetrate the marketplace, with three companies producing products -- Hitachi, Fujitsu and NEC. However, with the introduction of the Cray 2 and 3, and Cray XMP and YMP super computers, it appears that Cray's domination of the global market place will be secure, at least for the time being.³

What makes a super computer different?

The way super computers "crunch numbers" make them special. For example, most of us work mathematical problems using scalar processing. That is, we do arithmetic one step at a time, often getting confused if we try to do step three before completing step two. Super computers are good at scalar processing but only because they possess a great deal of sheer speed and power. Because everything in scalar processing must be done in a prescribed sequence, it is slow and cumbersome. It requires that each individual case be worked out horizontally, step by step.

On the other hand, super computers treat lists of numbers as vectors, performing repetitive calculations simultaneously, instead of working through a series of single calculations. A vector processor can work rapidly through these lists, performing all the multiplications at the same time. This adds enormously to the actual operating speed especially for applications programs which include a large percentage of repetitive operations, like "do

loops" in FORTRAN.

The exceptional performance of super computers is in large part due to their ability to do vector processing. This unique ability, coupled with writing programs that optimize the number of operations that can be performed, results in significant speed improvements. Since scientific problems tend to involve many repetitive calculations, "vectorizing" programs also increases operational speed.

Problems run on super computers tend to be much more complicated than getting the right pay checks out to thousands of employees. Through efficient vector processing (hardware) and skillful vectorization (software), solutions can be obtained quickly -- often in seconds and minutes, or in a few hours.⁴

Significant improvements in super computing architectures to the practical limits of the laws of physics are not expected to result in more than a 20 fold improvement in super computer performance. Therefore, new architectural approaches are required to achieve the performance desired by today's scientists, such as massively parallel processors, will be discussed later in this paper.

Super computing - The Government's Role to Date

The U.S. government's demand for faster and faster computer operating speeds has provided a stable market for super computers, and was the real genesis of the field of super computing. The method of procurement used by government to develop and build super computers, had a definite affect on those companies that were eventually successful in developing commercially viable super computers.

This analysis is based largely upon the analysis of Lee Holcomb of the Massachusetts Institute of Technology (MIT).⁵

Holcomb distinguished between two types of government - industry relationships, that of "friendly buyer" or that of "fly off."

"In the friendly buyer relationship, the government agrees to purchase supercomputers from industry prior to their development. In such a relationship the government sets very loose "requirements" and provides considerable software development support after delivery. By contrast, in the fly off relationship, the government contracts to purchase supercomputers only after they pass detailed system specifications. Most fly off acquisitions involve two or more contractors in the early phase of development. A competition is held and the government selects one firm for further development phases, with all other firms being eliminated from that round of development and acquisition."⁶

The end result of these two approaches to development and acquisition has been, quite frankly, a mixed bag. Sperry and Control Data Corporation were the first "developers" of high speed computers for the U.S. government and are no longer in this segment of the industry (Sperry has merged into a new corporation known as UNISYS, with Burroughs). IBM, also involved in early development work, is presently only half-heartedly involved and is most recently fighting to keep market share and regain profitability, following its first year of profit loss in 1991. While loss of profitability cannot be attributed to development and procurement methodologies of super computers, an examination of the corporate survivors can be.

In the mid-sixties, CDC stepped into the super computing niche left by Sperry and IBM, with its announcement of the 6600. It was the most powerful computer available and proved to be the first

commercially successful super computer, eventually delivering 59 during its production. The 6600 had been developed by a small group of engineers, lead by Seymour Cray, who is still recognized as the foremost designer of super computers. The CDC supercomputer was a third generation computer using integrated circuits, with discrete components connected by shortened pieces of wire. Innovative packaging resulted in exceptional performance that was five times that of earlier computers, for about half the price.⁷

The government procured the CDC 6600 as a "friendly buyer". It had no specific performance requirements but accepted CDC's specifications and agreed that the Atomic Energy Commission's Lawrence Livermore National Laboratories (LLNL) would be the first user. CDC wanted the prestige of a national laboratory behind the 6600 to help in its commercial ventures. Despite penalty payments for late delivery, the 6600 became a classic and established CDC's reputation as a manufacturer of high performance scientific computers.⁸

CDC followed the 6600 with the 6800 and the 7600, the first two production models being delivered to LLNL in 1969. These machines were procured using the "friendly buyer" approach. However, the next requirement for super computing was the National Aeronautics and Space Administration (NASA), using a "fly off" procurement. The requirement was for a parallel processor using 64 processing elements at the University of Illinois. Burroughs Corporation was selected competitively to build the machine known as ILLIAC IV. The machine ended up costing 5 times more than

estimated and never achieved desired computational speeds. Only one was ever produced and Burroughs abandoned super computers.⁹

In the meantime, CDC began a "fly off" relationship with the government, again for LLNL. In 1964, CDC was selected to build an experimental vector processor for a fixed price contract of \$24 million. Technical and management problems led to delays, increased development cost and late delivery, with penalty payments. CDC absorbed several million dollars in development costs to deliver only four STAR 100 (SString Array Processor) computers. Because of the requirement to "vectorize" software, the computer was not popular with programmers, but after software conversion, applications ran as much as five times faster than earlier CDC computers.

In 1972, Seymour Cray left CDC to form Cray Research. CDC went on to build the Cyber 203/205 series of computers using lessons learned during the development of the STAR. This was also a successful line of super computers, and one in which the government was a "friendly buyer", taking delivery of early models and funding significant software development.¹⁰

The government was also a "friendly buyer" for the Cray-1, which was produced for less than \$5 million between 1972 and 1976. This was the computer Seymour Cray left CDC to build. The government has been a major user and contributor to software for the Cray machines, significantly enhancing their hardware's commercial viability. Clearly, the "friendly buyer" approach has not only benefitted the government, but Cray as a super computer

manufacturers as well.

Recently LLNL backed out of a contract to purchase the first Cray 3 super computer from Cray Computer Corporation (Seymour Cray's new start-up company) after the company missed a critical December 1991 deadline. Instead the Department of Energy will buy a C90 super computer from Cray Research Incorporated. The Cray 3 was designed by Seymour Cray, who left his namesake company in 1989 to establish Cray Computer to build the machine. The Cray 3 is more than two years behind schedule and was considered too risky by Cray Research. Both computers are based upon a 16 processor architecture and use leading edge gallium arsenide technology. Gallium arsenide is several times faster than silicon and more tolerant of changes in temperature.¹¹ Both the C90 and Cray 3 cost about \$30 million and will achieve peak processing speeds of 16 gigaflops. The impact of government's decision to cancel this contract with Cray's new startup company, remains to be seen.¹²

The bottom line is, the US government has played a significant role in determining the players in the top-end of the computer industry. Despite reduced budgets for defense and other areas of government expenditures, the government can be expected to continue to play an important and highly influential role in shaping this segment of the computer industry. It is however, incumbent upon government to measure previous successes and failures to determine its future "leadership" role.

Leading Edge Computing

"The bottom line is that any country which seeks to control its future must effectively exploit high performance computing. A country which aspires to military leadership must dominate, if not control, high performance computing. A country seeking economic strength in the information age must lead in the development and application of high performance computing in industry and research."

White House Science Committee - Nov 1985

As discussed previously, Cray Research remains the country's only manufacturer that is commercially successful in selling super computers. Despite this fact, it is misleading to think that the United States is down to one manufacturer to face the bulk of the competition from Japan and the European Community. For example, IBM is trying to reenter the super computer market after a 20 year hiatus. New companies have begun development of new types of super computers. Convex Computer Company took the approach of developing "minisupercomputers" that rival the performance of the Cray computer of only a few years earlier. Increasingly, these manufacturers of high performance computers are becoming more successful domestically and internationally.

Other highly experimental designs are under development or have been prototyped by manufacturers. Thinking Machines Corporation is developing a prototype computer containing 65,536 processors. These "massively parallel" designs face significant hurdles before they become commercially viable, but they serve a

major function of stretching the state of the art. This often referred to, "lunatic fringe" of computing, is financed by venture capitalists and is frequently run by recently graduated hardware and software engineers turned entrepreneur. These extremely risky ventures are frequently the source of significant technological advancement. The purpose of this section is to look at two leading hardware developers and several critical technical areas that will eventually result in significant improvement in processing speed and power.

Thinking Machine Corp.

In November 1991, the Connection Machine 5, manufactured by Thinking Machine Corp, became the fastest super computer in the world, bettering the most powerful Cray computers by a factor of 100. Coupled with this benchmarking success, seven customers have purchased the CM-5, paying from \$1.5 million to \$25 million for models that contain from 32 to 1,024 processors.

The success of the Connection Machine 5 marks several milestones in computer science. The first and most important, is commercial acceptance of "massively parallel" computers as a viable architecture. Presently, there are estimated to be more than a half dozen start-up companies selling or developing parallel-processing computer architectures. Digital Equipment Corp and IBM, the two largest computer manufacturers, have endorsed the concept and are working on designs (IBM formed a joint venture with Thinking Machines in September 1991).

Danny Hillis, designer of CM-5, said his modular design is capable of being configured from 32 to 16,000 processors and could deliver a peak speed of 2 teraflops -- 2 billion scientific calculations per second. Given today's size and prices for components, the machine would fill a small gymnasium and cost approximately \$200 million. Currently this computer is too large and expensive. Analysts predict that a \$50 million commercially viable small size teraflop machine, will be available by the mid-90's.

Japan has also selected "massively parallel" computers as one of their technology targets for long-term research. At least three Japanese manufacturers -- NEC, Hitachi and Fujitsu -- are busy making Connection Machine-like computers of their own.

Intel Corporation

The Intel Corporation is generally associated with its central processing units and associated silicon chip manufacturing expertise. While it also manufactures computers, it is not generally known for top-end computer manufacturing. Yet, Intel's three year Touchstone Program is one of the country's most promising massively parallel computing projects. Touchstone is a jointly funded research and development project. Intel is funding \$19.9 million, with the Defense Advanced Research Projects Agency (DARPA) providing \$7.6 million.

The program is working toward established milestones -- IOTA, GAMMA and DELTA prototype systems -- with a goal of producing the

fourth, called SIGMA, by the end of 1992. SIGMA will be an attempt at manufacturing the fastest machine in the world and will consist of 2000+ processors, capable of operating at 150 billion floating point operations per second (gigaflops). Despite reported speed differences highlighted in the earlier section of this paper, between Thinking Machine and Intel super computers, there is a more important difference.

Intel's iPSC (Intel parallel super computer) designers are using multiple instruction, multiple data (MIMD) architectures, meaning they break a program into several discrete programs that run on separate parts of the machine concurrently. Each part or node is connected in a "hypercube" configuration in which each node is connected to a set number of neighboring nodes. On the other side, Thinking Machine is the main proponent of the single instruction, multiple data (SIMD) approach, in which data is spread out among the parallel nodes (up to 64,000 in the case of the Connection Machine). Data are then processed simultaneously, one instruction at a time.

Moving "leading edge" technology quickly into production has been one of the main thrusts of Intel's designers. The IOTA stage produced a concurrent in/out facility that the company incorporated into its iPSC 2 machines. GAMMA converted the 80386-based processing nodes to i860-based processors and was the direct predecessor of the iPSC-860 line. The DELTA stage will include up to 500 processing nodes and will offer performance up to 32 gigaflops. DELTA will also move to a new interconnection scheme

between nodes. While Intel has no current plans to market the DELTA machine, it has sold one to a recently formed super computer consortium, Concurrent Supercomputing Consortium, of 14 research institutions that will share the DELTA. One of the major objectives of the consortium, will be to identify strengths and weaknesses of the MIMD and SIMD approaches, as well as address the development of software compilers, that will optimize software for either approach.

The Intel and DARPA SIGMA machine will incorporate new Intel designed chips and be used as a testbed for new super computing technology.¹³

"Critical Technologies"

The speed of super computers will increase only by about a factor of two using currently materials (assuming architecture remains constant), which is insufficient to solve the types of scientific problems the applied research and development communities want to solve. While some researchers are turning to massively parallel architectures to increase speeds, like Thinking Machines and Intel, researchers are also looking at expanding the super computing performance envelope through advanced materials. The purpose of this section is to briefly examine some of these technologies to gain an appreciation of the research taking place that will lead to significant enhancements in performance. The importance of this research cannot be underestimated, as some researchers are predicting that this research will eventually allow

super computers to operate at petaflop calculation speeds (1×10^{15}), or 1000 X's faster than the currently sought teraflop machine.

I have for simplicity, divided these critical technical research areas into four categories;

- communications
- optics
- superconductivity
- heat dissipation

Communications. Communication bottlenecks are one of the major problems facing designers. Whether the problem is internal or external to the super computer, the large amount of data being processed between various components, slows down due to data path constraints as it is being passed between major super computer components. In addition to the various hardware and software design schemes (hypercube vice mesh node interconnect, or MIMD vice SIMD), new materials and production techniques will allow for significant increases in transmission speeds. Research at LLNL is showing progress in using high-temperature superconducting materials to create data transmission lines among computer components. These transmission lines can be made one-tenth the size of conventional metal, down to about 10 microns distance between the center of one line and the next. This will allow for either ten times the number of connections or reducing the size of the component. This method not only allows for shorter

connections, but facilitates faster transmission.

Optics. Optics, along with superconductors, is also critical to reducing communication bottlenecks. New semiconductor materials have increased efficiency up to 50%, are high power, no larger than a fingernail and will be coming down in price by a factor of 10. Newly developed solid state lasers are now compact, rugged and, most importantly, becoming cheap. Advances, coupled with improvements in high temperature superconductor components, will significantly increase super computer speeds.

Superconducting. At the component level, two important areas of research are showing promise. Researchers at Sandia National Laboratories have constructed a working transistor out of superconducting materials. This is the first active electronic device made from high-temperature superconductors. At the same time, Varian has constructed transistors out of diamond, that may have a larger impact than those made of superconductors. Diamond is radiation hardened, has excellent thermal conductivity and high electron mobility, resulting in very high speed switching. As you might imagine, cost of diamond transistors are currently prohibitive and production problems are unresolved.

Heat Dissipation. Another area being researched is component cooling. "Microchannel cooling" is a new process that uses etching techniques to cut a narrow deep channel into silicon substrates through which liquid coolants can circulate. Using water flow through the microchannels, the system can dissipate 50 times the amount of heat of current systems.

The bottom line is, researchers are pursuing a number of independent technical research projects. The result will be significantly higher speeds, reduced noise, low electrical losses and higher efficiency . . . that will eventually result in a super computer that is capable of achieving petaflop calculation speeds.¹⁴

The Super Computing Market Place

Despite the failures of many of the earlier manufacturers of super computers, the use of scientific and technical computing is growing significantly. The super computer marketplace is projected to grow by 21.7% per year through 1993, with parallel-processing machines growing by 64.4% annually. Compare this with the personal computer and mid-range commercial computer market, which are predicted to increase by 12.6% and 6.1% annually, you begin to get an appreciation for the growth of this segment of computing.

Super computers range in price from \$14,000 to \$23 million and are the most expensive computers to own and operate. Despite the high cost and operating expense, there are a number of companies with products available. The following is a fairly inclusive list of supercomputer manufacturers in the United States operating in 1991.

Supercomputer Suppliers¹⁵

<u>Company</u>	<u>Description</u>
Alliant Computer Systems Corp Littleton, Mass	

HPCCI & Industrial Policy

FX Series	Moderately parallel mini-supercomputer
FX/2800	Moderately parallel mini-supercomputer
BBN Advanced Computers Inc. Cambridge, Mass	
TC2000	Massively parallel
Concurrent Computer Corp. Trinton Fall, NJ	
5000 Series	Moderately parallel
6000 Series	Moderately parallel
8000 Series	Moderately parallel RISC based real-time
Convex Computer Corp Richardson, Texas	
C Series	Moderately parallel
Cray Research Inc. Minneapolis, MN	
XMS	Entry supercomputer
X-MP EA Series	Moderately parallel
X-MP	Moderately parallel
Encore Computer Corp. Fort Lauderdale, Fla	
Multimax	Symmetrical Multi- processor, mini
90 Family	Real-time symmetrical multiprocessor
Floating Point Systems Inc. Beaverton. Ore	
System 500	SPARC-based super- computer for SUN workstation/networks

HPCCI & Industrial Policy

FPS 5800	Array processor - attached to host
Fujitsu America Inc. San Jose, Ca	
VP-2000	Vector processing supercomputer
Harris Corp. Computer Systems Div Fort Lauderdale, Fla	
Nighthawk Series	Moderately parallel mini-supercomputer
Hitachi Data Systems Santa Clara, Ca	
S-800 Series	Large-scale super- computer
HNSX Supercomputers Inc Burlington, Mass	
SX-X	Moderately parallel
SX-2	Supercomputer
SX-3	Moderately parallel supercomputer
Intel Corp. Supercomputer Division Beaverton, Or	
iPSC/860	Massively parallel
MasPar Computer Corp. Sunnyvale, Calif	
MP-1	Massively parallel
nCube Beaverton, Or	
nCUBE 2	Massively parallel
Sequent Computer Systems Inc. Beaverton, Or	

HPCCI & Industrial Policy

Symmetry Series	General purpose parallel processing
Star Technologies Inc. Sterling, Va	
VP Series	Array processors - boost performance of host computer
Star 910	SUN SPARC-compatible server for network
Think Machine Inc. Cambridge, Mass	
Connection Machine System	Massively parallel

As mentioned earlier, massively parallel super computers are the fastest growing segment of the computer marketplace as more and more applications become available to take advantage of the architecture. The biggest challenge remaining, will be for these systems to win users over to new software development techniques (vectorization, and MIMD or SIMD) required to produce highly parallelized code that exploits the power of massively parallel machines.

The seventeen manufacturers listed in the table above, account for about 80% of current high-end super computer systems available. Additionally, Amdahl, Control Data Corp and IBM can provide vector processors to supplement their mainframes (basically the same approach used by the Japanese to reach super computer speeds), providing super computer performance.¹⁶

GOVERNMENT'S SUPERCOMPUTER INITIATIVES

As little as ten years ago, the United States had virtually no competition in manufacturing the fastest and most powerful computers. This abruptly came to an end in the summer of 1982, when Fujitsu announced it would market a computer six to eight times more powerful than the fastest American machines. This was followed in 1983 by an announcement that the Japanese had computers faster than the best Cray and CDC super computers.¹⁷

Japanese Initiatives

Another computing research project that has attracted attention is the "Fifth Generation" computer, started by Japan under the sponsorship of the Ministry of International Trade and Industry (MITI) in 1982. The objective of this research is to develop a new type of computer, superseding the previous four generations that were largely based upon vacuum tubes, transistors, integrated circuits and very large scale integrated circuits. The emphasis of this research is to create computers able to use human languages and perform like the human brain, or at least to act as a transparent link or collaborator, between the human operator and the problem being worked.¹⁸ The Fifth Generation project is often confused with the Japanese high speed computer project, which was initiated about the same time. This project is essentially different from the relatively straight forward "number crunching"

requirements of super computing, and is outside the scope of this paper.

Typical with most Japanese industrial initiatives, the Supercomputer Project involves a partnership between the Ministry of International Trade and Industry (MITI) and six major electronic manufacturers -- Fujitsu, Hitachi, Mitsubishi, Nippon Electric Company, Oki, and Toshiba. The initial project began in 1982 and was scheduled to run through 1989 with total funding projected at \$200 million.¹⁹ Consistent with the Japanese approach, research and development projects were divided between company and government laboratories, and information shared between potential commercial competitors.

The three most successful Japanese manufacturers of super computers are Fujitsu, Hitachi and Nippon Electric Company (NEC). In late 1983 Fujitsu and Hitachi delivered super computer systems to the Universities of Nagoya and Tokyo respectively. Both manufacturers' products were of similar design and were compatible with existing IBM mainframe software.²⁰ All three Japanese manufacturers have marketing agreements with American companies.

The Japanese high speed computer project has been successful at both the component (chip) and system (computer) levels. In 1990, benchmarking (measures how quickly a computer solves a batch of problems) of Japanese super computers against the fastest Cray, demonstrated 50% more computational speed. Both NEC and Fujitsu demonstrated speed advantages over the fastest Cray²¹.

In the last several years, the Japanese have demonstrated a

near obsession with super computing. A big advantage: Japanese suppliers make their own microchips, which are generally faster than those that Cray uses. A problem that the Japanese have continued to have difficulty cracking, is software. Sophisticated world markets demand the best software and Cray has over 500 programs available for their machines. Recently, Japan has chosen the Unix operating system as the system of choice and will begin standardizing their software programs. This could result in significantly narrowing the software gap that now exists.²²

The most recent information available on this project indicates that the Japanese are specifically looking at international teaming arrangements to address technical areas in which they need help. It also appears that U. S. developers are significantly ahead in "massively parallel" architectures and the software required to support this computing environment (discussed earlier in this paper).

European Supercomputers

While in the past, the national and international programs in Europe to advance information processing technology were denigrated, the advent of the European Community may change all that. Initially, France and Germany were the only countries in Europe that announced intentions to develop and manufacture super computers. A cooperative program known as the European Strategic Programme for Research and Development (ESPRIT I and, as of last year, II) is underway that specifically addresses earlier

programmatic problems. ESPRIT focuses on five technical areas; advanced microelectronics, software technology, advanced information processing, office automation and computer integrated manufacturing.²³ While current programs have not met with tremendous success, the European program possesses tremendous potential given its scientific capability, which is comparable with the US and Japan, coupled with a growing need for processing power.

Problems have begun to surface within EC '92, for continued funding for projects like ESPRIT. After years of pouring billions in subsidies into European electronics companies, the EC is calling for cuts. The EC has spent some \$6.2 billion in high-tech research since 1987. While it provided employment, it did little to boost Europe's competitive standing. With EC backing out of the subsidy and protection business, much of the debate will migrate to Europe's capitals. The outcome is not clear. What is clear though, is that Europe received little return on its investment. It appears that European manufacturers are looking to the U. S. and Japan, for teaming arrangements to obtain needed technologies.²⁴

Although European initiatives merit monitoring, they are not believed to be a threat to super computer developments or expected to become a major competitor for U.S. super computer manufacturers.

US SUPERCOMPUTER INITIATIVES

On July 10, 1990, the U. S. Senate approved the High Performance Computing Act (HPCA)²⁵. The key component in this

legislation was the establishment of a Federal Interagency High-Performance Computing Task Force, headed by the Secretary of Energy and Department of Energy to maximize efficient utilization of assets and ensure that results are shared within the research and development community, government, industry and academia. This represents a significant change in government policy, which had a number of relatively uncoordinated efforts funded with each doing their own thing.

The Task Force is composed of the Secretary of Energy, Secretary of Commerce, Secretary of Defense, Administrator of the National Aeronautics and Space Administration, and Director of the National Science Foundation. This group is responsible for the development of a comprehensive, long-range plan that sets goals for research, development and application of super computing. The program is required by law to address the following;

- Create a National Research and Education Network (NREN) capable of multi-gigabit-per-second data transmission capacity. The network should be operational by the end of 1996.
- Establish Collaborative Consortia that will address hardware and software research and development. Each consortia will be lead by a national laboratory, with participants from industry, Federal laboratories, and educational institutions.
- Solicit advice of potential users throughout development.
- Transfer of technology to the industrial sector is to be of the highest priority.

Estimated cost to the Federal Government (HPCA 1991):

(by fiscal year, in millions of dollars)

	1991	1992	1993	1994	1995	
Authorized	65	100	135	170	205	26

The impact of the HPCA could be significant. The National Science Foundation estimates that results could be an increase in productivity of American researchers by up to 200 percent. Such an increase in productivity with NREN and networked super computers would significantly increase benefits from the \$76.6 billion spent per year by the Federal Government on research and development.

In January 1992, the Bush Administration proposed a broad reaching plan, as part of the Fiscal 1993 Federal budget, significantly increasing research and development expenditures, specifically targeting areas that could increase productivity and improve national economic performance²⁷. The budget proposes that funding for High Performance Computing and Communications be increased by \$148 million, or 23 percent, to a total of \$803 million for Fiscal Year 1993. The goal is to assist in the development of a computing capability 1,000 times faster and a communications system 100 times faster than those in current use by 1996. In addition to increased resources, there are several other aspects of the program;

- Accelerate the pace of technology transfer by increasing the number of Cooperative Research and Development Agreements
- Expand role of the National Laboratories. Labs are expected to play an increasingly

important role in helping to form R&D consortia and other collaborative arrangements led by industry and universities.

- The budget proposes to make the Research and Experimentation tax credit permanent.
- Encourage R&D by Multinational Companies. Proposes an 18-month extension in rules governing foreign and domestic expenditures for companies with foreign operations.

Investments in research, technology and development are to be focused in four High-Performance Computing and Communications program areas;

- Hardware - Development of technology necessary to sustain parallel computer operations of trillions of operations per second on large projects. With efficient technology transfer, these systems will be the foundation of new commercial super computers and workstations.
- Advanced Software - Development of generic software to realize performance potential of highly parallel and networked super computers.
- National Research and Education Network - a significant upgrade of existing federally supported networks to provide a distributed computing capability. This gigabit network will become the foundation for sophisticated commercial networks.
- Basic Research and Human Resources - Support long-term research, increasing the number of students in computer science and transfer technology for industrial grand challenge applications.

In the first year of operation, 1991, HPCCI (nee HPCA) was successful in a number of technical and programmatic areas. Major new high performance computing systems were either announced or delivered. New software applications were developed for emerging high performance systems. Traffic on data networks doubled as the

HPCCI & Industrial Policy

number of interconnected local and regional networks increased. Federal agencies have begun solicitation to fund high performance computing research groups, centers and consortia. In addition, a large number of researchers, scholars, students, scientists, and engineers have been trained to use these new technologies.

The early involvement of industry and academia in the HPCCI program has already resulted in several critical recommendations;

- Expand the vision of HPCCI
- Establish a policy foundation for the future
- Improve management and increase opportunity for industry participation
- Reorder budget to balance program, especially in the area of software development

These recommendations are being addressed, with an additional \$68 million being provided for software development in the 1993 budget²⁸. The highly interactive and dynamic approach currently called for in the HPCCI program is already having results, and most importantly, government, industry and academia are beginning to establish channels of communication, albeit not optimally.

CONCLUSION AND RECOMMENDATIONS

The information covered thus far, projects a reasonably positive picture. While the HPCCI certainly has significant and far reaching potential, as a paradigm for industrial policy, it is important to digress and place the initiative in the context of current political and economic reality. Most economists agree that the last decade and a half has seen dramatic decreases in industrial output and productivity in many of the world's industrial economies. Inflation and unemployment have been higher than at any time in the post-World War II era. During this period, the United States has experimented with several economic theories in an attempt to fix problems.

The first attempt was to apply the theory of "supply-side economics". This was based upon the assumption that lower taxes would stimulate economic growth, and the resulting economic growth would offset a loss of government income through higher taxable income. In theory, this would more than off-set short term losses. However, this economic policy resulted in record deficits, ultimately causing the Federal government to "crowd out" industry and consumers in global capital markets to service deficits. Now, in the 1990's, there is increasing interest in the development of national industrial policy, largely due to the perceived success of Japan in strategic economic planning.

This second theory is based upon the observation that long-term government policy can result in significant and sustained economic growth. The theory goes something like this . . . the Japanese Ministry of International Trade and Industry (MITI), working in close concert with industry and banks, develops strategies that facilitate Japanese penetration of world automobile markets, development of automated steel mills, manufacture of dynamic random access memory chips and now production of the ultimate super computer.²⁹

Two leading advocates of national industrial policy, Magaziner and Reich in their book Minding America's Business suggest, "that US companies and the government develop a coherent and coordinated industrial policy whose aim is to raise the real income of our citizens by improving the patterns of our investments rather than by focusing only on aggregate investment levels"³⁰, like HPCCI. On the other hand, Schultze, an economist with the Brookings Institute, argues that this approach, by necessity, includes the picking of industrial "winners" and "losers"³¹. He also argues, that the American system of government is incapable of picking and choosing between firms or regions to arrive at efficient, productive and high value-added industries. These arguments represent the extremes of the industrial policy debate. Long-term planning, with maximum involvement of government, industry and academia targeting critical technical areas, does not necessarily lead to a priori, selection of "winners" or "losers". The High Performance Computing Act of 1991, as modified by the HPCCI in the

1993 Federal Budget, represents a reasonable alternative, while providing a potential paradigm for a national industrial policy, coupled with an implemented strategy.

The significance of the High Performance Computing Act, captured and renamed in the President's Budget for Fiscal Year 1993, is that the United States has put in place a policy specifically aimed at increasing the overall productivity of American industry, by targeting critical technical areas. The impact of improving productivity can be dramatic. For example, if productivity were to grow in the 90s at the same rate it did during the 60's, the average American workers salary would increase by 30 percent.³² Currently, the American worker's purchasing power has been in steady decline since the 1960's. The convergence of computer and communication technology, coupled with the ability to effectively integrate both technologies from corporate boardroom, through engineering and production design, to the factory floor offers tremendous potential for manufacturing productivity increases. Along the same lines, American education, medical and social service delivery, could also be expected to improve, providing these sectors were targeted. They are currently being considered for inclusion in HPCCI.

Significant return on research and development investment through increased productivity, is also supported by studies done by the Bureau of Labor Statistics. In a 1989 study of productivity, the bureau estimates that the direct influence of R&D on productivity growth was greatest in manufacturing, accounting

for an average annual productivity growth of 0.49 percent between 1948 - 1967³³. Simply, our historical "ad hoc" approach to long-term government investment can be improved upon and without picking losers or winners.

Based upon a two year study of key U. S. industries, the Council on Competitiveness, a group of businesses and universities, observed that American firms do best in technologies requiring low capital investment. Particularly those areas of research and development undertaken by individuals or spawned from basic research conducted in business and government labs. By contrast, America falls behind in technologies that require massive upfront funds and long development lead times, that are often the result of "targeting" by foreign governments or joint research ventures among overseas firms.³⁴

The super or high performance computing and communications initiatives, coupled with government policies surrounding their development could be seminal to achieving the greatest increases in industrial productivity since mass production. The high performance or super computer paradigm is significant, and assuming a long-term commitment by American political leadership to invest in these technologies, could provide the engine for another industrial revolution based upon "information management", in its broadest sense.

Despite changes in government policy and industrial strategy, there are additional recommendations that should be addressed;

HPCCI & Industrial Policy

-> Establish a cabinet level position and department for the Chief Science Advisor and Office of Science and Technology Policy, or evolve the recently established Critical Technologies Institute. This not only raises the visibility of research and development, but will assist in information exchange and integration. According to recent information provided during visits with industry, the exchange of meaningful information and cooperation with government leadership, scientists and laboratories, remains the greatest impediment to real progress. While the actual research would be performed by appropriate government laboratories and their departments, the visibility and integration of critical technologies is integral to productivity enhancements. Additionally, the scope of application should be widened to include broad-based societal applications, such as medical care or education delivery.

-> Government must make a long-term commitment to industrial strategy that targets critical technologies. This commitment, coupled with a national strategic vision, is critical to security and economic prosperity. Consortia need to know that government's commitment is long-term and relatively immune to periodic shifts in political winds. First in line, would be restructuring the current 25% R & D tax credit, which currently hampers long-term research and planning rather than stimulating them. Currently, Congress must vote to re-enact it every year. Moreover, the tax credit applies only to increases in R & D, which tempts firms to

halt research and restart it sometime in the future. This tax credit should be made permanent.

-> Emphasis must be on "end results", not the organizations and their bureaucracies. "Flat" or horizontal organizations, focused on specific technical targets must be unencumbered and free to fail, without the threat of funds withdrawal. Management will be key to establishment of creative and innovative high-risk environments. This will require a significant cultural change for government.

-> Ensure that information is readily available to the community performing research. NREN will provide a vehicle for information exchange. Regional consortia with government, industry and academia would provide research "critical mass", with important technology transfer to local industry.

-> Encourage competition between technical solutions. We have seen that MIMD and SIMD are two very different software approaches to utilizing massively parallel computers. Both hold promise for significant improvement in the speed at which super computers can process information. Neither is right or wrong and should be given full opportunity to reach their respective potential.

-> "Users", in the broadest sense, must be involved in

development seminally. Front-end involvement not only addresses training and educational issues, but ensures that products are targeted against real world problems as early as possible. Evolutionary development also get products into the marketplace soonest . . . resulting in a faster return on investment and increasing the scope of the initiative.

-> Government must go out of its way to eliminate its adversarial relationship with industry, whether real or perceived. This is critical to changing the nature of American industry..

-> The smooth and efficient transfer of technology to industry must be the bottom line of government industrial strategy. This is something that the U. S. has not done well and must become a major objective of a fully integrated R&D community.

The real test of this strategy will, once again, be the integration and leverage these technologies provide American industry. A single cabinet level department, will ultimately be needed to promulgate long-term industrial vision, policy and strategy, in close concert with industry and academia. This cooperative arrangement is critical to competition, enhanced productivity and ultimately profitability (read standard of living).

Clearly, national industrial policy, is only part of the equation. The government also needs to develop a coherent long-

term macroeconomic policy that will support capital formation and investment. The finest research in the world, will not result in better products or increased profits without an economic environment that provides incentives for industrial growth and investment.

James Madison clearly understood the equation for national power, when he wrote in Federalist Paper No. 41, "A system of government meant for duration ought to contemplate these revolutions and be able to accommodate itself to them." While he addressed the coming industrial revolution, he recognized that government must change to fit economic realities if it is to survive. During the same period, Frederick the Great, after the Seven Years War had devastated Prussia, supported an industrial policy of "ein Plus machen" or, simply, make a profit. In the intervening two hundred years, little has changed.

The ultimate strategy of the United States, necessary to realize security, peace and prosperity, resides with an energetic Hamiltonian executive, that effectively communicates vision and generates national will. The United States is at an important juncture, where leadership, especially in the form of vision, is vital. The balance of "power" and "profit" is also critical. In this regard, the development of a super computing development community with government, industry and academic involvement is an opportunity to forge a new environment.

This is not to imply that the industrial paradigm of the High Performance Computing and Communication Initiative represents a

panacea that will solve perceived industrial malaise, it is however a reasonable basis for the development and implementation of an industrial vision, policy and strategy, that targets critical technologies. Technologies that will be the eventual centerpieces for increasing productivity and profitability, creating new high value added industries, increasing employment and living standards, and ultimately increasing the security and national power of the United States. Strategic long-term vision from committed political leadership is vital, as is consistent and reinforcing industrial strategy to replace that which is fragmented and contradictory. Real problems deserve real answers. This paradigm, coupled with macroeconomic, and educational strategies, is a step in the right direction.

ENDNOTES

1. Kahn, Robert E., "The Quest -- A New Generation in Computing," IEEE Spectrum, The Institute of Electrical and Electronics Engineers, Inc., November 1983, p.36.
2. Holcomb, Lee B., "The U.S. Supercomputer Industry: A Strategy for the Future," Massachusetts Institute of Technology, Cambridge, MA, 1984, p. 17.
3. Karin, Sidney and Norris, Parker Smith, The Supercomputer Era, Harcourt Brace Jonavich, Orlando, 1987, p. 87-89.
4. Karin and Smith, p. 41 - 43.
5. Holcomb, p. 30.
6. Holcomb, p. 48.
7. Holcomb, "The U.S. Supercomputer Industry," pp. 49-50.
8. Ibid., p. 51.
9. Ibid. p. 50.
10. Ibid. p. 52.
11. Smale, John. Introduction to Telecommunications Systems, 1986, p.33.
12. Marsan, Carolyn D., "Livermore Swaps One Cray Firm For Another", Federal Computer Week, January 13, 1992, pp. 30-31.
13. Davis, Dwight D. "Supercomputers". Datamation. May 1, 1991, p.38.
14. Henderson, Breck W. "Exotic Materials to Boost Speed of Next Generation of Computers" Aviation Week and Space Technology. April 15, 1991. pp 66-67.
15. Faulkner Technical Reports, Pennsauken, NJ, 1991.
16. Abbott, Lawrence. "Supercomputers: Big Bang, Big Bucks", Datamation, February 15, 1991. pp 73-77.
17. Walsh, John, "Leadership in Computer Technology," Science, Vol 219, January 7, 1983, p.1.

HPCCI & Industrial Policy

18. Feigenbaum, Edward A. and McCorduck, Pamela, The Fifth Generation, New American Library, New York, p. 11.
19. "Supercompeting over Supercomputers," Science, Vol 220, May 6, 1983, p 583.
20. Fernbach, Sidney, "Supercomputers -- Past, Present, Prospects," Future Generations Computer Systems, North-Holland, Vol. 1, No. 1, July 1984, pp 23-30.
21. "Supercomputers: Guess Who's Streaking Past Cray", Business Week. December 10, 1990, p 212 - 216.
22. Ibid., p. 216.
23. "ESPRIT: Europe Challenges US and Japanese Competitors", Future Generations Computer Systems, North-Holland, Vol 1, No 1, July 1984, p 61-62.
24. "European High Tech Tries to do Something Drastic: Grow Up", Business Week., March 25, 1991, p. 48.
25. "Department of Energy High-Performance Computing Act of 1990", 101st Congress, 2d Session, Senate Report 101-377, pp 2-3.
26. Ibid., p. 24.
27. "The Budget for Fiscal Year 1993", Office of Management and Budget, The White House, Part One-87.
28. Ibid, Part One-101.
29. Schultze, Charles L., "Industrial Policy: A Dissent", The Brookings Review, Fall 1988, p 3.
30. Magaziner, Ira and Reich, Robert. Minding America's Business, Harcourt Brace Javanovich, New York, 1982, p. 4.
31. Schultze, p. 4.
32. Krugman, Paul. The Age of Diminished Expectations. The MIT Press. Cambridge, Mass. p. 174.
33. Ibid., Part One - 92.
34. Dentzer, Susan. "Staying Ahead in High Tech"., US News and World Report. April 1, 1991, p. 53.

BIBLIOGRAPHY

- Abbott, Lawrence. "Supercomputers: Big Bang, Big Bucks", Datamation, February 15, 1991, pp. 73-77.
- Davis, Dwight D. "Supercomputers", Datamation, May 1, 1991, p. 38.
- Dentzer, Susan. "Staying Ahead in High Tech", US News and World Report, April 1, 1991, p. 53.
- "Department of Energy High-Performance Computing Act of 1990", 101st Congress, 2d Session, Senate Report 101-377, pp. 2-3.
- "European High Tech Tries to do Something Drastic: Grow Up", Business Week, December 10, 1990, pp. 212 - 216.
- "ESPRIT: Europe Challenges US and Japanese Competitors", Future Generations Computer Systems, North-Holland, Vol 1, No 1, July 1984, p. 61-62.
- Faulkner Technical Reports, Pennsauken, NJ, 1991.
- Feigenbaum, Edward A. and McCorduck, Pamela, The Fifth Generation, New American Library, New York, p. 11.
- Fernbach, Sidney, "Supercomputers -- Past, Present, Prospects", Future Generations Computer Systems, North-Holland, Vol 1, No 1, July 1984, pp. 23-30.
- Henderson, Breck W., "Exotic materials to Boost Speed of Next Generation of Computers", Aviation Week and Space Technology, April 15, 1991, pp. 66-67.
- Holcomb, Lee B., "The US Supercomputer Industry: A Strategy for the Future", Massachusetts Institute of Technology, Cambridge, MA, 1984, p. 17.
- Kahn, Robert E., "The Quest -- A New Generation in Computing", IEEE Spectrum, The Institute of Electrical and Electronics Engineers, Inc., November 1983, p. 36.
- Karin, Sidney and Norris, Parker Smith, The Supercomputer Era, Harcourt Brace Jonavich, Orlando, 1987, pp. 87-89.
- Krugman, Paul, The Age of Diminished Expectations, The MIT Press. Cambridge, MA, p. 174.
- Magaziner, Ira and Reich, Robert, Minding America's Business, Harcourt Brace Jonavich, New York, 1982, p. 4.

Marsan, Carolyn D., "Livermore Swaps One Cray Firm for Another", Federal Computer Week, January 13, 1992, pp. 30-31.

Schultze, Charles L., "Industrial Policy: A Dissent", The Brookings Review, Fall 1988, p. 3.

Smale, John. Introduction to Telecommunications Systems, 1986, p. 33.

"Supercomputers: Guess Who's Streaking Past Cray:", Business Week, December 10, 1990, pp 212 - 216.

"Supercompeting over Supercomputers", Science, Vol 220, May 6, 1983, p. 583.

"The Budget for Fiscal Year 1993", Office of Management and Budget, The White House, Part One-87.

Walsh, John, "Leadership in Computer Technology", Science, Vol 219, January 7, 1983, p. 1.